

Learning Detailed Face Reconstruction from a Single Image

Elad Richardson¹

Matan Sela¹

Roy Or-El²

Ron Kimmel¹

¹Department of Computer Science, Technion - Israel Institute of Technology

²Department of Computer Science and Engineering, University of Washington

{eladrich, matansel, ron}@cs.technion.ac.il

royorel@cs.washington.edu



Figure 1: Results of the proposed network. Reconstructed geometries are shown next to the corresponding input images.

Abstract

Reconstructing the detailed geometric structure of a face from a given image is a key to many computer vision and graphics applications, such as motion capture and reenactment. The reconstruction task is challenging as human faces vary extensively when considering expressions, poses, textures, and intrinsic geometry. While many approaches tackle this complexity by using additional data to reconstruct the face of a single subject, extracting facial surface from a single image remains a difficult problem. As a result, single-image based methods can usually provide only a rough estimate of the facial geometry. In contrast, we propose to leverage the power of convolutional neural networks to produce a highly detailed face reconstruction from a single image. For this purpose, we introduce an end-to-end CNN framework which derives the shape in a coarse-to-fine fashion. The proposed architecture is composed of two main blocks, a network that recovers the coarse facial geometry (CoarseNet), followed by a CNN that refines the facial features of that geometry (FineNet). The proposed networks are connected by a novel layer which renders a depth image given a mesh in 3D. Unlike object recognition and detection problems, there are no suitable datasets for training CNNs to perform face geometry reconstruction. Therefore, our training regime begins with a supervised phase, based on synthetic images, followed by an unsupervised phase that uses only unconstrained facial images. The accuracy and robustness of the proposed model is demonstrated by both qualitative and quantitative evaluation tests.

1. Introduction

Faces, with all their complexities and vast number of degrees of freedom, allow us to communicate and express ourselves through expressions, mimics, and gestures. Facial muscles enable us to express our emotions and feelings, while facial geometric features determine one's identity. However, the flexibility of these qualities make the recovery of facial geometry from a flat image a challenge. Moreover, additional ambiguities arise as the projection of a face onto an image depends also on its texture and material properties, lighting conditions, and viewing direction.

Various methods mitigate this uncertainty by using additional data such as a large photo collection of the same subject [36, 35, 21, 28, 33], continuous video frames [43, 39, 5, 11] or a rough depth map [43, 18]. In many cases, however, we only have access to a single facial image. In this setup, common schemes can be divided to 3D morphable model (3DMM) techniques [3, 4], template-based methods [20, 15] and data-driven approaches [26, 40, 34].

Here, we propose an end-to-end neural network for reconstructing a detailed facial surface in 3D from a single image. At the core of our method is the idea of breaking the reconstruction problem into two phases, each solved by a dedicated neural network architecture. First, we introduce CoarseNet, a network for recovering the coarse facial geometry as well as the pose of the face directly from the image. To train CoarseNet, a synthetic dataset of facial images with their matching face geometry and pose is synthetically generated. The rough facial geometries are modeled using a 3DMM [3], which provides a compact representation that can be recovered using the proposed network. However,

this representation can only capture coarse geometry reconstruction. Next, in order to capture fine details, we introduce FineNet, a network that operates on depth maps and thus is not constrained by the morphable model representation. FineNet receives a coarse depth map, alongside the original input images and applies a shape-from-shading like refinement, capturing the fine facial details. To train FineNet, we use an unlabeled set of facial images, where a dedicated loss criterion is introduced, to allow unsupervised training. Finally, to connect between the CoarseNet 3DMM output and the FineNet depth map input, we introduce a novel layer which takes the 3DMM representation and pose parameters from CoarseNet, and produces a depth map that can be fed into FineNet. This layer supports back-propagation to the 3DMM representation allowing joint training of the two networks, possibly refining the weights of CoarseNet.

The usage of an end-to-end network here is exciting as it connects the problem of face reconstruction to the rapidly expanding applications solved by CNNs, potentially allowing us to further improve our results following new advances in CNN architectures. Moreover, it allows fast reconstructions without the need for external initialization or post-processing algorithms. The potential of using a CNN for reconstructing face geometries was recently demonstrated in [34]. However, their network can only produce the coarse geometry, and must be given an aligned template model as initialization. These limitations force their solution to depend on external algorithms for pose alignment and detail refinement.

The main contributions of the proposed method include:

- An end-to-end network-based solution for facial surface reconstruction from a single image, capable of producing detailed geometric structures.
- A novel rendering layer, allowing back-propagation from a rendered depth map to the 3DMM model.
- A network for data refinement, using a dedicated loss criterion, motivated by axiomatic shape-from-shading objectives.
- A training scheme that bypasses the need for manually labeled data by utilizing only synthetic data and unlabeled facial images.

2. Related Work

Automatic face reconstruction attracts a lot of attention in the computer vision and computer graphics research communities. The available solutions differ in their assumptions about the input data, the priors and the techniques they use. When dealing with geometry reconstruction from a single image, the problem is ill-posed. Still,

there are ways for handling the intrinsic ambiguities in geometry reconstruction from one image. These solutions can be roughly divided into the following categories:

3DMM Methods. In [3], Vetter and Blantz introduced the 3D Morphable Model (3DMM), a principal components analysis (PCA) basis for representing faces. One of the advantages of using the 3DMM is that the solution space is constrained to represent only likely solutions, thereby simplifying the problem. While the original paper assumes manual initialization, more recent efforts propose an automatic reconstruction process [4, 47]. Still, the automated initialization pipelines usually do not produce the same quality of reconstructions when only one image is used, as noted in [33]. In addition, the 3DMM solutions cannot extract fine details since they are not spanned by the principal components.

Template-Based Methods. An alternative approach is to solve the problem by deforming a template to match the input image. One notable paper is that of Kemelmacher-Shlizerman and Basri [20]. There, a reference model is aligned with the face image and a shape-from-shading (SfS) process is applied to mold the reference model to better match the image. Similarly, Hassner [15] proposed to jointly maximize the appearance and depth similarities between the input image and a template face using SIFT-flow [24]. While these methods do a better job in recovering the fine facial features, their capability to capture the global face structure is limited by the provided template initialization.

Data-Driven Methods. A different approach to the problem uses some form of regression to connect between the input image and the reconstruction representation. Some methods apply a regression model from a set of sparse landmarks [1, 10, 25], while others apply a regression on features derived from the image [22, 7]. [26] applies a joint optimization process that ties the sparse landmarks with the face geometry, recovering both. Recently, a network was proposed to directly reconstruct the geometry from the image [34], without using sparse information or explicit features. That paper demonstrated the potential of using a network for face reconstruction. Still, it required external procedures for fine details extraction as well as initial guess of the face location, size, and pose.

In a sense, the proposed solution combines all of these different procedures. Specifically, a **3DMM** is used to define the input for a **Template-Based** refinement step, where both parts are learned using a **Data-Driven** model.

3. Coarse Geometry Reconstruction

The first step in our framework is to extract the coarse facial geometry and pose from the given image. Our solution is motivated by two recent efforts, [34] which proposed to train a network for face reconstruction using synthetic data,

and [46] which solved the face alignment problem using a network. Although the methods focus on different problems, they both use an iterative framework which utilizes a 3D morphable model. The proposed method integrates both concepts into a holistic alignment and geometry reconstruction solution.

3.1. Modeling The Solution Space

In order to solve the reconstruction problem using a CNN, a representation of the solution space is required. To model the facial geometries we use a 3D morphable model [3], where an additional blendshape basis is used to model expressions, as suggested in [8]. This results in the following linear representation

$$S = \mu_S + A_{id}\alpha_{id} + A_{exp}\alpha_{exp}. \quad (1)$$

Where μ_S is the average 3D face, A_{id} is the principal component basis, A_{exp} is the blendshape basis, and α_{id} and α_{exp} are the corresponding coefficient vectors. A_{id} and A_{exp} are collected from the Bosphorus dataset [37] as in [34], where the identity is modeled using 200 coefficients, and the expression using 84.

For projecting the 3D model to the image plane, we assume a parallel weak perspective projection.

$$\begin{bmatrix} p_x \\ p_y \end{bmatrix} = \begin{bmatrix} f & 0 & 0 \\ 0 & f & 0 \end{bmatrix} [R|t] \begin{bmatrix} P_x \\ P_y \\ P_z \\ 1 \end{bmatrix}, \quad (2)$$

where p, P are the pixel location in the image plane and in the world coordinate system, respectively, f is the focal length, and $[R|t]$ is the extrinsic matrix of the camera. Hence, the face alignment is modeled using only 6 parameters: 3 Euler angles, a 2D translation vector and a scale. The pose parameters are normalized so that a zero vector would correspond to a centralized front facing face. Overall, we have a representation of 290 parameters for both geometry and pose. We will denote this representation as r .

3.2. The CoarseNet Training Framework

The realization that the power of single-pass systems is limited, has made the application of iterative networks popular. While some methods [38, 23] use a cascade of networks to refine their results, it has been shown that a single network can also be trained to iteratively correct its prediction. This is done by adding feedback channels to the network that represent the previous output of the network as a set of feature maps. The network is then trained to refine its prediction based on both the original input and the feedback channels. This idea was first proposed by Carreira *et al.* in [6].




Figure 2: Feedback Representation. (a,d) are masked input images, (b,e) are the corresponding PNCCs of the network's output and (c,f) are the resulting normal maps.

3.2.1 Feedback Representation

Defining the feedback channels of the previous output of the network is crucial, as it would affect the overall performance of our iterative framework. Roughly speaking, we would like the feedback channels to properly represent the current state of the coarse facial geometry. In practice, different types of feedback channels would emphasize different features of the current state. For instance, in [46] the Projected Normalized Coordinate Code (PNCC) was introduced. This feature map is computed by first normalizing the average face and painting the RGB channels of the current vertices with the x, y and z coordinates of the corresponding vertex on the average model, see Figures 2b and 2e.

Next, we propose to use the normal map as an additional channel, where each vertex is associated with its normal coordinates. These normal values are then rendered as RGB values. The purpose of the normal map is to represent more local features of the coarse geometry, which are not emphasized by the PNCC. The proposed solution uses both feedbacks, creating a richer representation of the shape. Examples of these representations are shown in Figure 2.

3.2.2 Acquiring the Data

In order to train the proposed framework, a large dataset of 3D faces is required. However, due to the complexity in acquiring accurate 3D scans for a large group of people, no such dataset is currently available. Note that unlike different annotations, such as landmark positions, which can be manually collected for an existing set of unlabeled images, the 3D geometry has to be captured jointly with the photometric data. A possible solution would be to apply existing reconstruction methods to 2D images, and use these reconstructions as labels. However, such an approach would limit the reconstruction quality to that of the reconstruction method we use.

Here, we choose to follow the line of thought proposed in [34] and create a synthetic dataset by drawing random representations of geometry and pose, r_{gt} , which are then rendered using random texture, lighting and reflectance. This process provides a dataset of 2D images, for which the pose and corresponding geometry are known by construction.

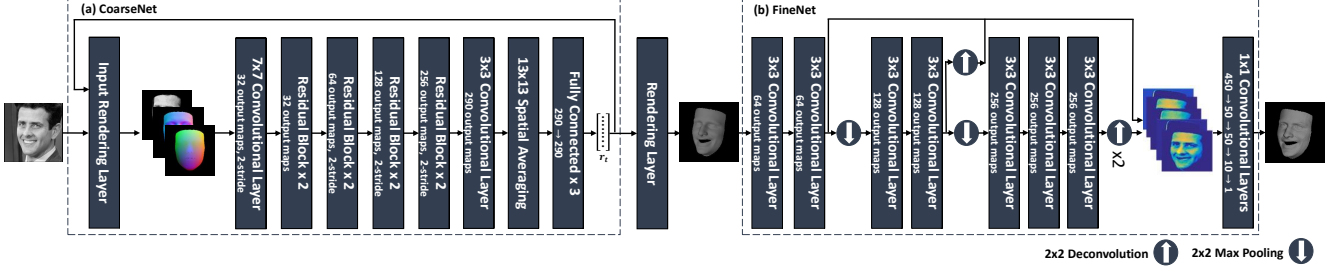


Figure 3: The End-to-End network, composed of CoarseNet, FineNet and the rendering layer.

The iterative refinement process is then simulated by drawing another set of parameters, r_t , which is sampled between r_{gt} and a random set of parameters, r_{rnd} .

$$r_t = \beta \cdot r_{gt} + (1 - \beta) \cdot r_{rnd}, \quad 0 \leq \beta \leq 1, \quad (3)$$

r_t represents the current estimation of the solution, and is used to generate the PNCC and normal map. The network is then trained to predict the ground-truth, r_{gt} , representation from the current one, r_t . Note, that unlike [34] our representation r captures not only the geometry, but also the pose. Hence r_{gt} and r_{rnd} can vary also in their position and orientation.

3.3. The CoarseNet Architecture and Criterion

CoarseNet is based on the ResNet architecture [17], and is detailed in Figure 3. Note that the input layer includes the feedback channel, and that a gray-scale image is used. The last element in the proposed architecture is the training criterion. As our representation is composed of both geometry and pose parameters, we choose to apply a different training criterion for each part of the representation. For the geometry we apply the Geometry Mean Square Error (GMSE) suggested in [34],

$$L(\hat{\alpha}, \alpha) = \|[A_{id}|A_{exp}] \hat{\alpha} - [A_{id}|A_{exp}] \alpha\|_2^2, \quad (4)$$

where $\hat{\alpha}$ is the geometry received from the network, and α is the known geometry. The idea behind GMSE is to take into account how the different coefficients affect the resulting geometry. For the pose parameters we found that a simple MSE loss over the 6 parameters is sufficient. We weigh the two loss criteria so that we get approximately the same initial error for both.

3.4. Using CoarseNet

We feed CoarseNet it with a 200×200 image of a face. Such an image can be automatically acquired using a standard face detector, such as the Viola-Jones detector [41]. The initial parameters vector, r_0 , is set to zeros, corresponding to a centered mean face μ_S . In addition, the input image is always masked in accordance with the visible vertices in

the feedback channel. The masking is applied in order to improve our generalization capability from synthetic data to real-world images, as our synthetic data is more accurate for the head region. Although the mask is inaccurate in the first iteration, it is gradually refined. The network is then applied iteratively, producing the updated geometry r_t , which is used to create the new feedback input. This process is repeated until convergence, as shown in Figure 4.

4. The Coarse to Fine Approach

For many tasks, such as face frontalization [47, 16], reconstructing the coarse geometry is sufficient. However, reconstructing fine geometric structures such as wrinkles could be useful for other applications. It is clear that while working in the morphable model domain, we cannot capture such details. To solve that, we transfer the problem to the unconstrained image plane, representing the geometry as a depth map. The role of the proposed FineNet would then be to modify the given coarse depth map, based on the original image, for capturing the fine details.

4.1. The Rendering Layer

To connect CoarseNet with FineNet we propose a novel rendering layer. The layer receives the geometry and pose representation vector as the input and outputs a 200×200




Figure 4: Progress through iterations. For each iteration the following are shown from top to bottom: The cropped input image, the PNCC and the normal map.




Figure 5: Gradient Flow. Gradients from FineNet are first propagated to the depth map, and then propagated from each pixel to the matching vertices. The gradients on the triangular mesh are then propagated back to the morphable model representation.

depth map of the geometry in the corresponding pose. This is done in two steps, first the 3D mesh is calculated from the geometry parameters and positioned above the image plane,

$$\begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix} = \begin{bmatrix} f & 0 & 0 \\ 0 & f & 0 \\ 0 & 0 & 1 \end{bmatrix} R [A_{\text{id}} | A_{\text{exp}}] \hat{\alpha} + \begin{bmatrix} t_x \\ t_y \\ 0 \end{bmatrix}. \quad (5)$$

The 3D mesh is then rendered using a z-buffer renderer, where each pixel is associated with a single triangular face from the mesh. In order to handle potential occlusions, When a single pixel resides in more than one triangle, the one that is closest to the image plane is chosen. The value of each pixel is determined by interpolating the z-values of the mesh face using barycentric coordinates

$$\tilde{z} = \lambda_0 z_0 + \lambda_1 z_1 + \lambda_2 z_2, \quad (6)$$

where z_i is the z-value of the i^{th} vertex in the respective triangle and λ_i is the corresponding coordinate. During back-propagation the gradients are passed from each pixel to the matching vertex, weighted by the corresponding coordinate,

$$\frac{dE}{dz_i} = \frac{dE}{d\tilde{z}} \frac{d\tilde{z}}{dz_i} = \frac{dE}{d\tilde{z}} \lambda_i, \quad (7)$$

where E is the loss criterion. Note that we assume that the barycentric coordinates are fixed. Alternatively, one could derive the coordinates with respect to x_i and y_i . In addition, no gradients are propagated to hidden vertices. A similar approach was applied for example in [48]. Finally, the gradients are propagated from each vertex back to the geometry basis, by taking the derivative of Equation 5 with respect to $\hat{\alpha}$. The gradients transfer is visualized in Figure 5.

4.2. FineNet Framework

Delicate facial features such as wrinkles and dimples are difficult to represent by a 3DMM low dimensional space, mainly due to their high diversity. Hence, in contrast to CoarseNet, we need to use a pixel-based framework to

recover the fine details. Recently, several notable pixel-based CNN architectures [12, 27, 14] were used for various fine grained tasks like semantic and instance segmentation [27, 14], optical flow [9], and human pose estimation [42]. First successful attempts to reconstruct surface normals using these architectures [2, 44] have motivated our FineNet architecture. The proposed framework differs from both these networks in its output (depth map vs. normal map) and training regime (unsupervised vs. supervised).

The FineNet is based on the hypercolumn architecture suggested in [14]. The main idea behind this architecture is to generate a per-pixel feature map which incorporates both structural and semantic data. This is achieved by concatenating the output responses from several convolution layers along the path of the network. Due to pooling layers, the output maps size of inner layers does not match the size of the input image, therefore, they are interpolated back to the original size, to create a dense per-pixel volume of features. This volume is then processed by several 1×1 convolution layers to create the final prediction.

We choose the VGG-Face [31] as a base for our hypercolumn network since it was fine tuned on a domain of faces. For interpolating, we apply a slightly different scheme than that of [14]. Instead of directly up-sampling each feature map to the original size using bilinear interpolation, we use cascaded 2-strided 2×2 up-convolution layers to up-sample the feature maps. This is done in order to improve the quality of the features, as the interpolation is now also part of the learning process. In contrast to recognition problems, refining the facial features is a relatively local problem. Therefore, we truncate the VGG-Face network before the third pooling layer and form a $200 \times 200 \times 450$ hypercolumn feature volume. This volume is then processed by a set of 1×1 convolutional layers used as a linear regressor. Note, that this fully convolutional framework allows us to use any size of input images. Figure 3 describes the FineNet architecture.

4.3. FineNet Unsupervised Criterion

To train FineNet some form of loss function is required. One possible solution would be to simply use an MSE criterion between the network output and a high-quality ground-truth depth map. This would allow the network to implicitly learn how to reconstruct detailed faces from a single image. Unfortunately, as mentioned in Section 3.2.2, a large dataset of detailed facial geometries with their corresponding 2D images, is currently unavailable. Furthermore, a synthetic dataset for this task cannot be generated using morphable models as there is no known model that captures the diversity of fine facial details. Instead, we propose an unsupervised learning process where the loss criterion is determined by an axiomatic model. To achieve that, we need to find a measure that relates the output depth map to the 2D image.

To that end, we resort to Shape from Shading (SfS).

Recent results in SfS [20, 45, 13, 30, 29] have shown that when given an initial rough surface, subtle geometry details can be accurately recovered under various lighting conditions and multiple surface albedos. This is achieved by optimizing some objective function which ties the geometry to the input image. In our case, an initial surface is produced by CoarseNet and its depth map representation is fed into FineNet along with the input image. We then formulate an unsupervised loss criterion based on the SfS objective function, transforming the problem from an online optimization problem to a regression one.

4.3.1 From SfS Objective to Unsupervised Loss

Our unsupervised loss criterion was formulated in the spirit of [30, 29]. The core of our loss function is an image formation term, which describes the connection between the network’s output depth map and its input image. This term drives the network to learn fine detail recovery and is defined as

$$E_{sh} = \left\| \rho \left\langle \vec{l}, \vec{Y}(\hat{z}) \right\rangle - I \right\|_2^2. \quad (8)$$

Here, \hat{z} is the reconstructed depth map, I is the input intensity image, ρ is the albedo image, and \vec{l} are the first-order spherical harmonics coefficients. $Y(\hat{z})$ represents the matching spherical harmonics basis,

$$Y(\hat{z}) = (1, n_x(\hat{z}), n_y(\hat{z}), n_z(\hat{z})), \quad (9)$$

where $(n_x(\hat{z}), n_y(\hat{z}), n_z(\hat{z}))$ is the normal expressed as a function of the depth. Notice that while I is an input to FineNet, the scene lighting \vec{l} and albedo map ρ are unknowns. Generally, the need to recover both lighting and albedo is part of the ambiguity in SfS problems. However, here we can utilize the fact we do not solve a general SfS problem, but one constrained to human faces. This is done by limiting the space of possible albedos to a low dimensional 3DMM texture subspace.

$$\rho \approx T = \mu_T + A_T \alpha_T. \quad (10)$$

where μ_T is the average face texture, A_T is a principal component basis and α_T is the corresponding coefficients vector. In our implementation, 10 coefficients were used.

Now, as shown in [20], the global lighting can be correctly recovered by assuming the average facial albedo, $\hat{\rho} = \mu_T$, using the coarse depth map, z_0 , as follows

$$\vec{l}^* = \underset{\vec{l}}{\operatorname{argmin}} \left\| \hat{\rho} \left\langle \vec{l}, \vec{Y}(z_0) \right\rangle - I \right\|_2^2. \quad (11)$$

Note that this is an overdetermined linear problem that can be easily solved using least squares. Given the lighting coefficients, the albedo can also be easily recovered as

$$\alpha_T^* = \underset{\alpha_T}{\operatorname{argmin}} \left\| (\mu_T + A_T \alpha_T) \left\langle \vec{l}^*, \vec{Y}(z_0) \right\rangle - I \right\|_2^2. \quad (12)$$



Figure 6: Light and albedo recovery. Images are presented next to the recovered albedo, rendered with the recovered lighting.

As in Equation 11, this is an overdetermined linear problem that can be solved directly. Based on the resulting albedo and lighting coefficients we can calculate E_{sh} and its gradient with respect to \hat{z} . A few recovery samples are presented in Figure 6. To regularize the solution, a fidelity and smoothness terms are added to the criterion of FineNet.

$$\begin{aligned} E_f &= \|\hat{z} - z_0\|_2^2, \\ E_{sm} &= \|\Delta \hat{z}\|_1, \end{aligned} \quad (13)$$

where Δ is the discrete Laplacian operator. These terms guarantee that the solution would be smooth and won’t stray from the prediction of CoarseNet. The final per-pixel loss function is then defined as

$$L(\hat{z}, z_0, I) = \lambda_{sh} E_{sh}(\hat{z}, I) + \lambda_f E_f(\hat{z}, z_0) + \lambda_{sm} E_{sm}(\hat{z}). \quad (14)$$

Where the λ s determine the balance between the terms and were set as $\lambda_{sh} = 1$, $\lambda_f = 5e^{-3}$, $\lambda_{sm} = 1$. The gradient of L with respect to \hat{z} is then calculated and used for backpropagation.

4.3.2 Unsupervised Loss - A Discussion

The usage of unsupervised criterion has some desired traits. First, it eliminates the need for an annotated dataset. Second, it ensures that the network is not limited by the performance of any algorithm or the quality of the dataset. This results from the fact that the loss function is entirely dependent on the input, in contrast to supervised learning SfS schemes such as [44] and [2], where the data is generated by either photometric stereo or raw Kinect scans, respectively. In addition, unlike traditional SfS algorithms, the fact that the albedo and lighting coefficients are calculated only as part of the loss function means that at test time the network can produce accurate results directly from the intensity and depth inputs, without explicitly calculating the albedo and lighting information. Although the CoarseNet can be trained to generate the lighting and albedo parameters, we chose not to include them in the pipeline for two reasons. First, the lighting and albedo are only needed for the training stage and have no use during testing. Second, both (11) and (12) are over-determined systems which can be solved efficiently with least squares, thus, using a CNN for this task would be an overkill.



Figure 7: Criterion Flow. Gradients from both loss criteria are propagated back to CoarseNet.

4.4. End-to-End Network Training

Finally, in order to train FineNet we connect it to CoarseNet using the proposed rendering layer which is added between the two networks. Thus, a single end-to-end network is created. We then use images from the VGG face dataset [31], and propagate them to the framework. The forward pass can be divided into three main steps. First, each such image is propagated through CoarseNet for four iterations, creating the coarse geometry representation, this can be done similarly to the forward pass in a Recurrent Neural Network. Then, the rendering layer transforms the 3DMM representation to a depth map. Finally, the depth map, alongside the original input image, is propagated through FineNet resulting in the dense updated depth map. The criterion presented in 4.3 is then used to calculate the loss gradient. The gradient is backpropagated through the network allowing us to train FineNet and fine-tune CoarseNet. Note that the fact that CoarseNet was already trained is crucial for a successful training. This stems from the fact that the unsupervised loss function depends on the coarse initialization, which cannot be achieved without the synthetic data.

In a sense, this is similar to the training process of transfer learning in Neural Networks, where a new network is concatenated to a pre-trained network in order to learn a new domain. This scheme fine tunes the pre-trained network to give the optimal inputs to the new network in order to learn the new task. In order to prevent CoarseNet from deviating too much from the original coarse solution, a fidelity criterion is added to CoarseNet’s output. This criterion is the MSE between the current CoarseNet solution and the original one. Gradients from both FineNet and the fidelity loss are then weighted and passed through CoarseNet, fine-tuning it, as presented in Figure 7.

5. Experiments

In order to evaluate the proposed framework we performed several experiments to test its accuracy on both 3D facial datasets and *in the wild* inputs. Both qualitative and quantitative evaluations are used to demonstrate the strength of the proposed solution. Our method is compared to the template based method of [20], to the 3DMM based method introduced as part of [47] and to the data driven method of [34]. Note that unlike our method, all of the above require alignment information. We use the state-of-the-art

alignment method of [19] to provide input for these algorithms.

For a qualitative analysis we show our results on 400×400 *in-the-wild* images of faces. As can be seen in Figure 10, our method exposes the fine facial details as opposed to [47, 34] and is more robust to expressions and different poses than [20]. In addition, we compare our reconstructions with a state of the art method for reconstruction from multiple images [36]. The results are shown in Figure 8, one can see that our method is able to produce a comparable high quality geometry from only a single image. Finally, Figure 9 shows our method robustness to different poses, while Figure 1 shows some more reconstruction results.

For a quantitative analysis of our results we used the Face Recognition Grand Challenge dataset V2 [32]. This dataset consists of roughly two thousand color facial images aligned with ground truth depth of each pixel. Each method provided an estimated depth image and a binary mask representing the valid pixels. For the purpose of fair judgment, we evaluated the accuracy of each method on pixels which were denoted as valid by all the methods. As shown in Table 1, our method produce the lowest depth error among the tested methods.

Finally, as noted in Section 4.2 the fully convolutional FineNet can receive inputs with varying sizes. This size invariance is a vital property for our detail extraction network, as it allows the network to extract more details when a high quality input image is available. Figure 11 shows that although our network was trained only on 200×200 images it gracefully scales up for 400×400 inputs.




Figure 8: (a) and (c) are two input images, (b) and (d) are their 3D reconstruction via the proposed method. (e) is a reconstruction of the same subject, based on 100 different images recovered with the method proposed in [36].



Figure 9: Method Robustness. Our method shows robustness to different orientations, even in nearly 90° angles.

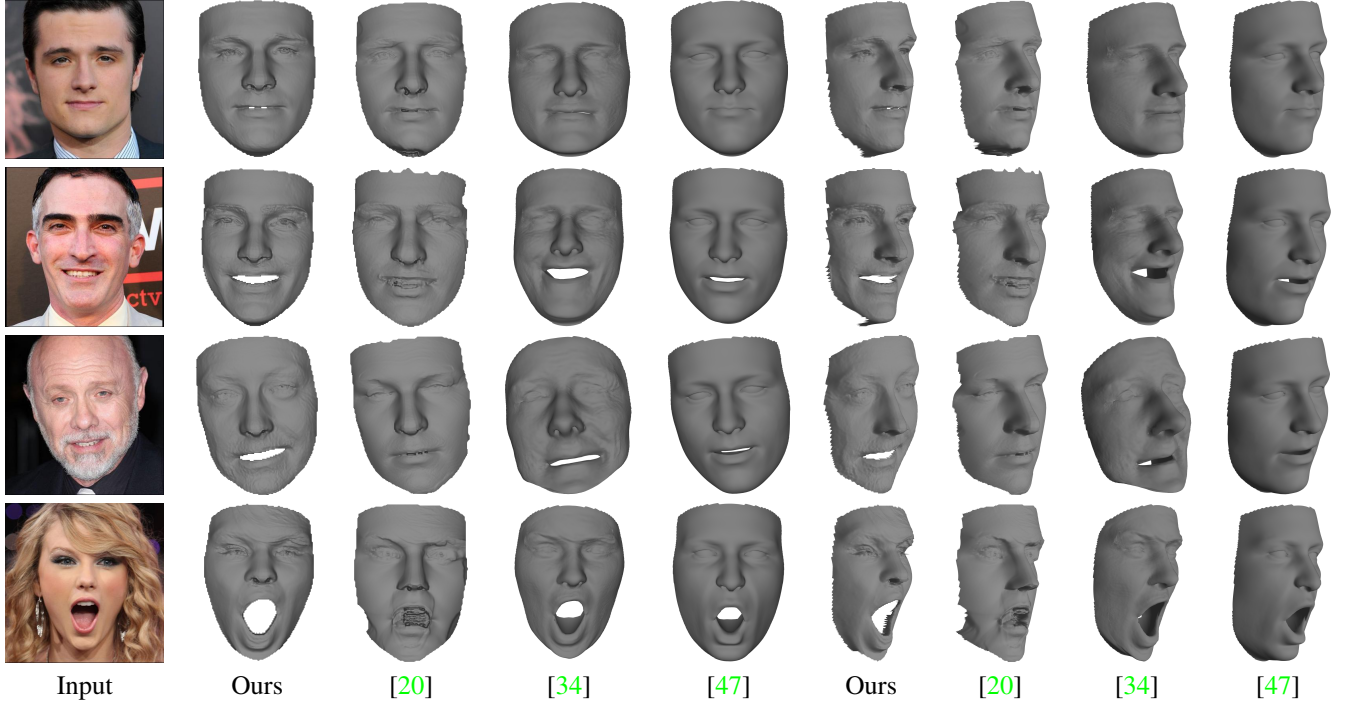


Figure 10: Qualitative Results. Input images are presented alongside the reconstruction results of different methods from two different viewpoints. Note that unlike the other methods, the proposed approach is robust to pose and expression variations, while still capturing subtle facial details.




Figure 11: Input Scaling. (a) is the input image and (b) is the coarse depth map from CoarseNet. In (c) the output of FineNet for a 200x200 input is presented, while in (d) a 400x400 input is used.

Method	Ave. Depth Err. [mm]	90% Depth Err. [mm]
Ours	3.22	6.69
[20]	3.33	7.02
[34]	4.11	8.70
[47]	3.46	7.36

Table 1: Quantitative Comparison. Depth estimation errors of the different methods are presented.

6. Conclusions

We proposed an end-to-end approach for detailed face reconstruction from a single image. The method is comprised of two main blocks, a network for recovering a rough estimation of the face geometry followed by a fine details reconstruction network. While the former is trained with synthetic images, the latter is trained with real facial images in an unsupervised training scheme. To allow unsupervised training, we introduced a rendering layer which enables the gradients to back-propagate through the entire network. As demonstrated by our comparisons, the proposed framework outperforms recent state-of-the-art approaches.

References

- [1] O. Aldrian and W. A. Smith. A linear approach of 3D face shape and texture recovery using a 3d morphable model. In *Proceedings of the British Machine Vision Conference*, pages 75–1, 2010. 2
- [2] A. Bansal, B. Russell, and A. Gupta. Marr revisited: 2D-3D alignment via surface normal prediction. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5965–5974, 2016. 5, 6
- [3] V. Blanz and T. Vetter. A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–

194. ACM Press/Addison-Wesley Publishing Co., 1999. 1, 2, 3
- [4] P. Breuer, K.-I. Kim, W. Kienzle, B. Scholkopf, and V. Blanz. Automatic 3D face reconstruction from single images or video. In *Automatic Face & Gesture Recognition, 2008. FG'08. 8th IEEE International Conference on*, pages 1–8. IEEE, 2008. 1, 2
- [5] C. Cao, D. Bradley, K. Zhou, and T. Beeler. Real-time high-fidelity facial performance capture. *ACM Transactions on Graphics (TOG)*, 34(4):46, 2015. 1
- [6] J. Carreira, P. Agrawal, K. Fragkiadaki, and J. Malik. Human pose estimation with iterative error feedback. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 3
- [7] M. Castela and J. Van Horebeek. 3D face shape approximation from intensities using partial least squares. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*, pages 1–8. IEEE, 2008. 2
- [8] B. Chu, S. Romdhani, and L. Chen. 3d-aided face recognition robust to expression and pose variations. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1907–1914. IEEE, 2014. 3
- [9] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 2758–2766, December 2015. 5
- [10] P. Dou, Y. Wu, S. K. Shah, and I. A. Kakadiaris. Robust 3D face shape reconstruction from single images via two-fold coupled structure learning. In *Proc. British Machine Vision Conference*, pages 1–13, 2014. 2
- [11] P. Garrido, M. Zollhöfer, D. Casas, L. Valgaerts, K. Varanasi, P. Pérez, and C. Theobalt. Reconstruction of personalized 3D face rigs from monocular video. *ACM Transactions on Graphics (TOG)*, 35(3):28, 2016. 1
- [12] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014. 5
- [13] Y. Han, J.-Y. Lee, and I. So Kweon. High quality shape from a single RGB-D image under uncalibrated natural illumination. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1617–1624, 2013. 6
- [14] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 447–456, 2015. 5
- [15] T. Hassner. Viewing real-world faces in 3d. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3607–3614, 2013. 1, 2
- [16] T. Hassner, S. Harel, E. Paz, and R. Enbar. Effective face frontalization in unconstrained images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4295–4304, 2015. 4
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 4
- [18] V. Kazemi, C. Keskin, J. Taylor, P. Kohli, and S. Izadi. Real-time face reconstruction from a single depth image. In *2014 2nd International Conference on 3D Vision*, volume 1, pages 369–376. IEEE, 2014. 1
- [19] V. Kazemi and J. Sullivan. One millisecond face alignment with an ensemble of regression trees. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 7
- [20] I. Kemelmacher-Shlizerman and R. Basri. 3D face reconstruction from a single image using a single reference face shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(2):394–405, 2011. 1, 2, 6, 7, 8
- [21] I. Kemelmacher-Shlizerman and S. M. Seitz. Face reconstruction in the wild. In *2011 International Conference on Computer Vision*, pages 1746–1753. IEEE, 2011. 1
- [22] Z. Lei, Q. Bai, R. He, and S. Z. Li. Face shape recovery from a single image using cca mapping between tensor spaces. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–7. IEEE, 2008. 2
- [23] H. Li, Z. Lin, X. Shen, J. Brandt, and G. Hua. A convolutional neural network cascade for face detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5325–5334, 2015. 3
- [24] C. Liu, J. Yuen, A. Torralba, J. Sivic, and W. T. Freeman. Sift flow: Dense correspondence across different scenes. In *European conference on computer vision*, pages 28–42. Springer, 2008. 2
- [25] F. Liu, D. Zeng, J. Li, and Q. Zhao. Cascaded regressor based 3D face reconstruction from a single arbitrary view image. *arXiv preprint arXiv:1509.06161*, 2015. 2
- [26] F. Liu, D. Zeng, Q. Zhao, and X. Liu. Joint face alignment and 3D face reconstruction. In *Proc. European Conference on Computer Vision*, Amsterdam, The Netherlands, October 2016. 1, 2
- [27] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. 5
- [28] F. Maninchedda, C. Häne, M. R. Oswald, and M. Pollefeys. Face reconstruction on mobile devices using a height map shape model and fast regularization. In *3D Vision (3DV), 2016 International Conference on*, pages 489–498. IEEE, 2016. 1
- [29] R. Or-El, R. Hershkovitz, A. Wetzler, G. Rosman, A. M. Bruckstein, and R. Kimmel. Real-time depth refinement for specular objects. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4378–4386, 2016. 6
- [30] R. Or-El, G. Rosman, A. Wetzler, R. Kimmel, and A. M. Bruckstein. RGBD-Fusion: Real-time high precision depth recovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5407–5416, 2015. 6
- [31] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *British Machine Vision Conference*, pages 41.1–41.12, 2015. 5, 7
- [32] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview

- of the face recognition grand challenge. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 947–954. IEEE, 2005. 7
- [33] M. Pietraschke and V. Blanz. Automated 3D face reconstruction from multiple images using quality measures. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3418–3427, 2016. 1, 2
- [34] E. Richardson, M. Sela, and R. Kimmel. 3D face reconstruction by learning from synthetic data. In *3D Vision (3DV), 2016 International Conference on*, pages 460–469. IEEE, 2016. 1, 2, 3, 4, 7, 8
- [35] J. Roth, Y. Tong, and X. Liu. Unconstrained 3D face reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2606–2615, 2015. 1
- [36] J. Roth, Y. Tong, and X. Liu. Adaptive 3D face reconstruction from unconstrained photo collections. *CVPR*, 2016. 1, 7
- [37] A. Savran, N. Alyüz, H. Dibeklioglu, O. Çeliktutan, B. Gökberk, B. Sankur, and L. Akarun. Bosphorus database for 3D face analysis. In *European Workshop on Biometrics and Identity Management*, pages 47–56. Springer, 2008. 3
- [38] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3476–3483, 2013. 3
- [39] S. Suwajanakorn, I. Kemelmacher-Shlizerman, and S. M. Seitz. Total moving face reconstruction. In *European Conference on Computer Vision*, pages 796–812. Springer, 2014. 1
- [40] S. Tulyakov and N. Sebe. Regressing a 3D face shape from a single image. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 3748–3755. IEEE, 2015. 1
- [41] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages I–511. IEEE, 2001. 4
- [42] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4732, June 2016. 5
- [43] T. Weise, H. Li, L. Van Gool, and M. Pauly. Face/off: Live facial puppetry. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '09*, pages 7–16, New York, NY, USA, 2009. ACM. 1
- [44] Y. Yoon, G. Choe, N. Kim, J.-Y. Lee, and I. S. Kweon. Fine-scale surface normal estimation using a single NIR image. In *European Conference on Computer Vision*, pages 486–500, 2016. 5, 6
- [45] L.-F. Yu, S.-K. Yeung, Y.-W. Tai, and S. Lin. Shading-based shape refinement of RGB-D images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1415–1422, 2013. 6
- [46] X. Zhu, Z. Lei, X. Liu, H. Shi, and S. Z. Li. Face alignment across large poses: A 3D solution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2, 3
- [47] X. Zhu, Z. Lei, J. Yan, D. Yi, and S. Z. Li. High-fidelity pose and expression normalization for face recognition in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 787–796, 2015. 2, 4, 7, 8
- [48] J. Zienkiewicz, A. Davison, and S. Leutenegger. Real-time height map fusion using differentiable rendering. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016. 5